AD-A214 622

(4)

GL-TR-89-0116

ENVIRONMENTAL RESEARCH PAPERS, NO. 1030

Rule-Based Systems for Visibility Forecasts

ROSEMARY M. DYER
GERALD L. FREEMAN, CAPTAIN, USAF

14 April 1989

Approved for public release; distribution unlimited.

ATMOSPHERIC SCIENCES DIVISION PROJECT 6670

GEOPHYSICS LABORATORY

HANSCOM AFB, MA 01731

"This technical report has been reviewed and is approved for publication"

FOR THE COMMANDER

DONALD A. CHISHOLM, Chief
Atmospheric Prediction Branch

ROBERT A. McCLATCHEY, Director
Atmospheric Sciences Division

This document has been reviewed by the ESD Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS).

Qualified requestors may obtain additional copies from the Defense Technical Information Center. All others should apply to the National Technical Information Service.

If your address has changed, or if you wish to be removed from the mailing list, or if the addressee is no longer employed by your organization, please notify AFGL/DAA, Hanscom AFB, MA 01731-5000. This will assist us in maintaining a current mailing list.

Unclassified
SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION Unclassified	1b RESTRICTIVE MARKINGS
2a SECURITY CLASSIFICATION AUTHORITY	3 DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited
2b DECLASSIFICATION / DOWNGRADING SCHEDULE	
4 PERFORMING ORGANIZATION REPORT NUMBER(S) GL-TR-89-0116 ERP, No. 1030	5 MONITORING ORGANIZATION REPORT NUMBER(S)
6a NAME OF PERFORMING ORGANIZATION Geophysics Laboratory	6b OFFICE SYMBOL (If applicable) LYP
7a NAME OF MONITORING ORGANIZATION	
6c ADDRESS (City, State, and ZIP Code) Hanscom AFB Massachusetts 01731-5000	7b ADDRESS (City, State, and ZIP Code)
8a NAME OF FUNDING / SPONSORING ORGANIZATION	8b OFFICE SYMBOL (If applicable)
9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER	
8c ADDRESS (City, State, and ZIP Code)	

10 SOURCE OF FUNDING NUMBERS

PROGRAM ELEMENT NO	PROJECT NO	TASK NO	WORK UNIT ACCESSION NO
62101F	6670	10	22

11 TITLE (Include Security Classification)
(U) Rule-Based Systems for Visibility Forecasts

12 PERSONAL AUTHOR(S)
Dyer, Rosemary M., and Freeman, Gerald L., Captain, USAF

13a TYPE OF REPORT	13b TIME COVERED	14 DATE OF REPORT (Year, Month, Day)	15 PAGE COUNT
Scientific, Final	FROM ______ TO ______	1989 April 14	34

16 SUPPLEMENTARY NOTATION

17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Fog forecasting, AI applications, Expert systems, Weather prediction

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

Zeus, a prototype rule-based system for forecasting visibility, was developed and tested at three locations along the East Coast. Evaluation of this system and further studies at AFGL have led to the conclusion that, although the usefulness of such a system has been demonstrated, a more fundamental approach to the representation of the knowledge base will have to be taken if such systems are to be deployed Air Force-wide. A discussion of the structure of a meteorological knowledge-base and an outline of how a generic fog forecast system might be developed are presented in this report.

20 DISTRIBUTION / AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION Unclassified
22a NAME OF RESPONSIBLE INDIVIDUAL Rosemary M. Dyer	22b TELEPHONE (Include Area Code) (617) 377-2967 — 22c OFFICE SYMBOL AFGL/LYP

DD FORM 1473, 84 MAR — 83 APR edition may be used until exhausted. All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE
Unclassified

Acknowledgement

A special thanks is extended to the forecasters and management of Detachment 2, 3rd Weather Squadron, Air Weather Service (AWS), who operate the weather forecasting office at Seymour-Johnson Air Force Base, North Carolina. Specifically, we want to recognize Major Love, the Detachment Commander; Mr. McKemie, the Chief Forecaster; and MSgt Strunk, the Station Chief. Thanks are also due to Detachment 1, 3rd Weather Squadron, at Shaw Air Force Base, South Carolina, for providing us with observational data and forecaster worksheets. The enthusiasm and dedication of these individuals enabled this evaluation to be accomplished, and thereby have aided the evolution of artificial intelligence applications in meteorology. We also are grateful for reviews and comments by our colleagues in the Atmospheric Prediction Branch and for manuscript preparation by Mrs. K.A. Campana.

Contents

Illustrations

Tables

Rule-Based Systems for Visibility Forecasts

1. INTRODUCTION

Between December 1985 and December 1986, a prototype rule-based expert system for forecasting selected visibility categories was developed. The system was a proof-of-concept prototype, limited to the forecast of restrictions to visibility by advective or radiative fog. Rules were developed independently for three locations along the Atlantic seaboard shown in Figure 1 (Dover AFB, Delaware; Seymour-Johnson AFB, North Carolina; and Fort Bragg, North Carolina), and preliminary performance evaluations were made.[1,2] When the system was delivered to AFGL, further evaluations were conducted, including added evaluation at one of the bases (Seymour-Johnson). The emphasis in this second phase was to determine the potential usefulness of such a system in operational USAF weather forecast offices.

Our experience with Zeus (the name given by GEOMET to the fog forecast system) is the basis of this report. There was a logical sequence to our efforts, which has been

(Received for publication 3 April 1989)

1. Stunder, M., Dyer, R., and Koch, R. (1987) The use of an expert system in assisting forecasters in visibility predictions, *3rd Converence on Interactive and Processing Systems in Meteorology, Oceanography, and Hydrology*, pp. 5206-5207.
2. Stunder, M., Koch, R., Sletten, T., and Lee, S. (1987) *ZEUS: A Knowledge-Based Expert System that Assists in Predicting Visibility at Airbases.* AFGL-TR-87-0019, AD A184197.

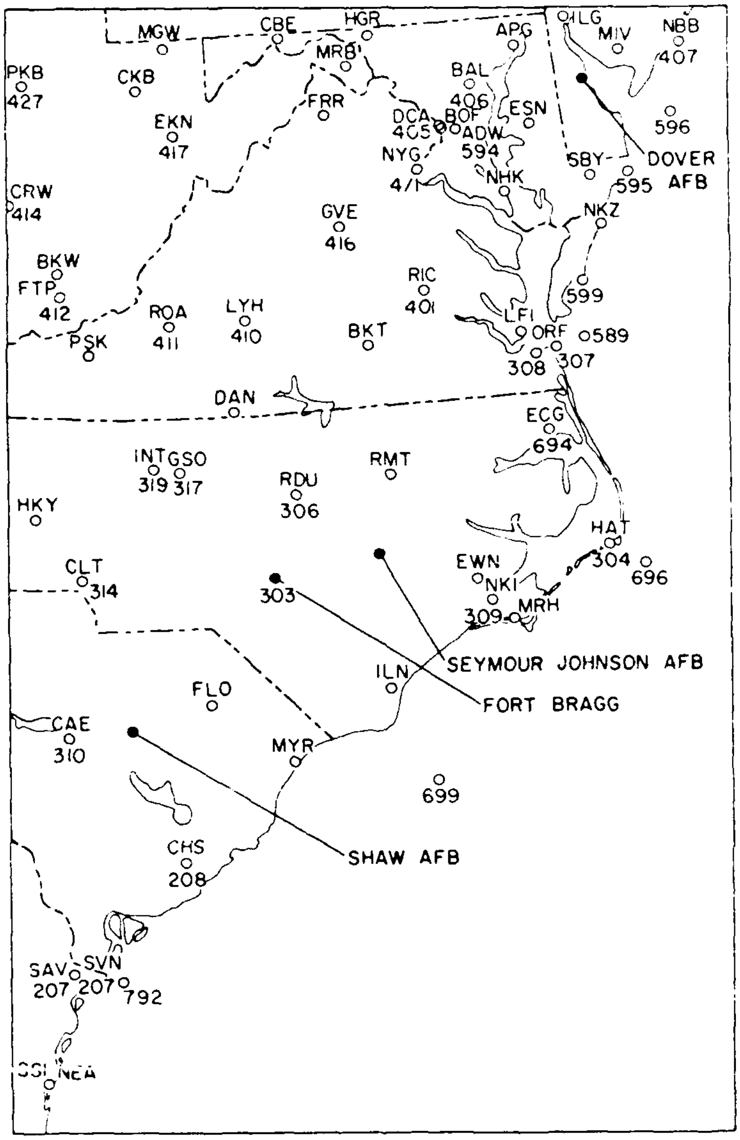

Figure 1. Locations of the Three Installations for Which Zeus Was Developed. The system developed for Seymour-Johnson was later adapted for Shaw AFB, also indicated on this map.

preserved in the arrangement of this report. Our first concern was the establishment of criteria for evaluating the system. It soon became apparent that simple performance evaluation was insufficient. The criteria for evaluation are discussed in Section 2 of this report. Some results of a thorough evaluation of the Seymour-Johnson system are given in Section 3. One of the criteria chosen is adaptability, which is defined as the ease with which an expert system designed for one geographic location can be used at another location, with little degradation in performance. As a test of Zeus's adaptability, the system designed

for Seymour-Johnson AFB was transferred to Shaw AFB, South Carolina, approximately 280 km southwest of Seymour-Johnson. Shaw is the station labeled "SSC" in Figure 1. The results are discussed in Section 4. Adapting the system to Shaw led to the concept of a knowledge-base structured in terms of general meteorological principles rather than site-specific rules of thumb. This "deep knowledge" approach to meteorological knowledge bases is discussed in Section 5. Finally, a summary of the lessons learned and recommenations for future applications of expert systems to short-range forecasts within the Air Force are given in Section 6.

2. EVALUATION CRITERIA

Evaluation of forecast techniques such as numerical models or radar products is a fairly straightforward statistical exercise. Here, objective measures of skill are computed using the new model or technique and the results are compared with a baseline procedure. The evaluation of expert systems in meteorology, however, is more complicated than the evaluation of techniques requiring less user interaction. We have identified five factors that should be included in the evaluation: verification, user acceptance, expandability, transportability, and adaptability These are explained below. Unfortunately, none of these factors is entirely independent of the others, and it is often difficult to assess the effect they have upon each other.

2.1 Verification

Verification, the comparison of the system's forecast with the actual weather at forecast time, is, of course, the bottom line in evaluating any forecast system or technique.* The difficulty in assessing the performance of an expert system, particularly in determining how the system can be improved, is that judgment calls by the user are required as part of the input data. Thus, verification is unalterably entwined with what we have termed "user acceptance." Nevertheless, an objective measure of performance must be established.

The measure chosen to compare the performance of human forecasters using the expert system with those not using the expert system is the Critical Success Index, or CSI.[3] This index was originally developed to compare techniques for forecasting severe storms, and eliminates the vast preponderance of occasions when the critical event did not occur

---

*In the artificial intelligence community, "verification" of expert systems is simply making sure that the knowledge entered into the system is correct and complete from the expert's point of view. What meteorologists call "verification," computer scientists call "validation." Throughout this report, we have used the meteorologists' definition of "verification."

3. Donaldson, R., Dyer, R., and Kraus, M. (1975) An evaluator of techniques for predicting severe weather events, *Ninth Conference on Severe Local Storms*, pp. 321-326.

and when nonoccurrence was correctly forecast. Ignoring these cases is not eliminating information. Rather, it allows us to concentrate on the significant differences between techniques, free of background noise. In this respect, it is analogous to the mathematical filtering of raw data before its spectrum analysis. Expert systems such as Zeus will be used more often by inexperienced forecasters, and generally in difficult or high stress situations. The CSI is a good measure of performance for those situations.

Another reason for choosing the CSI is its capability of incorporating any bias or tilt the forecaster might have toward forecasting the critical event (in this case, low visibility) when a strictly objective method would have predicted nonoccurrence. This tilt is present whenever the perceived penalty for failure to predict the phenomenon exceeds the penalty for a false alarm. It seems to be present in the fog predictions at the three bases studied by GEOMET. An advantage of using the CSI as a measure of performance is that the amount of tilt becomes explicit and, if desired, can be incorporated into the confidence factors associated with each visibility category in the expert system.

2.2 User Acceptance

The factor of user acceptance implies more than mere user-friendliness and ease of data input. In its current configuration, Zeus requires the input of judgment calls by the user exercising the system. The clarity of the questions, the appearance of a logical sequence to the questions, and the sensibleness of the information required all contribute both to the confidence of the user in the system and the accuracy of the prediction made by the system. Each accurate forecast, in turn, adds to the user's confidence. However, it should be noted that the inexperienced users--the ones most in need of the assistance afforded by an expert system--are the most likely to give inaccurate responses to the requests for judgment calls. Therefore, although there are no objective measures of user acceptance, an evaluation should be made based on user comments concerning specific requests for information, as well as their agreement or disagreement with the system's forecast.

2.3 Expandability

All expert systems are written in modular form. For example, Zeus has two main forecast modules--one for radiation fog, the other for advective fog. All such systems should be capable of expansion without extensive restructuring of the system's architecture. For example, it should be easily possible to add modules for the forecast of visibility deterioration because of precipitation, haze, or blowing snow or sand. Further, the structure of the expert system should be such that the addition of automatic data input from whatever source should be relatively straightforward. Finally, results should be easily incorporated into other expert systems.

2.4 Transportability

"Transportability" is a term used by computer scientists to indicate the possibility of systems developed on mainframe computers or specialized computers such as Symbolics being used on other computers such as the Z-100 found in operational weather stations throughout the Air Force.

2.5 Adaptability

By their very nature, expert systems designed to predict weather phenomena in the 0-to-6-hour time frame rely heavily upon the experience of local experts and rules of thumb tailored to the specific location of intended use. It is unavoidable, therefore, that some modification be required to adapt a system written for one locality to another location. If this modification cannot be made without repeating the entire knowledge acquisition process, then either the subject of the system is unsuited to a network of forecast stations, or the method of knowledge representation is too superficial. An expert system can be deemed adaptable if the rules can be modified without recourse to extensive knowledge acquisition and if the verification statistics for the modified system are not signifcantly poorer than those for the original system.

3. EVALUATION OF ZEUS AT SEYMOUR-JOHNSON AFB

As part of the original contract, GEOMET provided verification statistics for the three versions of its system, using archived data from the three Air Force installations as well as user evaluation at each forecast office. A more complete evaluation was then performed at GL, using the criteria described in Section 2 of this report. Many factors muddy the waters during the evaluation of expert systems--factors that are either totally absent or have minimal effect on conventional computer programs. Most of these involve synergism between user and computer, a relationship that constantly evolved during the course of the evaluation.

The method of data input changed during this period. At first, all input--both data and judgment calls--were entered manually in response to questions from the expert system. This was soon changed so that all temperatures, winds, cloud cover, and visibility used in any of the rules were entered before the program began. In this version of the program, locations of pertinent synoptic features were still entered manually, with the choices expressed in terms of specific geographic features. After some protests from some users, a third version of the program included a background map, with sectors indicating locations of interest. A further refinement, not available at the conclusion of the contract, will have current and predicted future positions of the synoptic features indicated by a mouse or a light pen.[1] These modifications to the user interface change the accuracy of the data input, and hence of the results.

Different users had different degrees of success in using the computer program. Many of the IF clauses in the rule base require judgment calls from the user. Will the low deepen during the next 4 hours? Will the high form a ridge to the southwest? The more experienced the user, the more likely he was to answer these questions correctly (and, ironically, the less need he had of a fog forecast expert system). In many cases, rerunning the program, but indicating how the synoptic situation actually changed rather than how the user forecast it to change, altered the expert system's forecast from an incorrect to a correct one. The effect of this can be demonstraed by looking at the verification statistics for Dover AFB. Real-time use of the system produced a skill score of 0.16. Using archived data from the same location, GEOMET obtained an average skill score of 0.55, with a range from 0.46 to 0.55 (Ref. 2, pp. 126, 135). At Seymour-Johnson, where users were apparently more familiar with computers and more accurate forecasters, the difference between real time and archived data was less dramatic. As time progressed and users became more familiar with the logic of Zeus, differences among users at Seymour-Johnson were masked by the factor discussed below.

After some experience with the system, users at Seymour-Johnson were able to manipulate results by responding in a way that would either force the computer to "kick out" of the system with an immediate forecast of continued good visibility (for those occasions when they were certain there would be no fog and they did not want to spend the time answering all the systems' questions), or prevent the computer from kicking out (for marginal occasions or for questions that they felt had a definite cutoff point). This obscured some deficiencies in the logic of the expert system--deficiencies that were uncovered only after painstaking examination of all possible combinations of true-false values for the IF clauses in all the rules. See Section 5 for further discussion of how systems can be made less fallible.

3.1 Verification

Skill scores and P-scores were computed by GEOMET for three test sites, both for real-time use by the operational forecaster on duty and for using archived data. The forecast was of visibility: When lower visibility is predicted because of a frontal passage or for any reason other than advective or radiation fog, the system gives a message to the effect that it was not designed to handle such situations. Visibility was expressed as one of three categories, the definition of which varies from base to base, depending on the type of aircraft used: Category 1 is less than 1 mile at Seymour-Johnson, less than 3/4 mile at Fort Bragg, and less than 1/2 mile at Dover; Category 2 is between 1 and 3 miles at Seymour-Johnson, between 3/4 and 2 miles at Fort Bragg, and between 1/2 and 2 miles at Dover; and Category 3 is greater than 3 miles at Seymour-Johnson and greater than 2 miles at Fort Bragg and Dover. Closer examination of the statistics by GL personnel led to a revaluation

of the performance ratings, this time using the critical success index (CSI) as the relevant verification statistic.

The CSI is applicable to the forecast of difficult-to-predict events, and it ignores those occasions on which the nonoccurrence was both predicted and experienced. These occasions comprise the bulk of the data and often mask differences in the performance of two techniques for predicting the rare event. The CSI is described in Donaldson et al.[3] Figure 2 illustrates the factors involved. X is the number of times the event was predicted and occurred; Y is the number of failures to detect when the event occurred but was not predicted; and the false alarms are represented by Z, the number of times the event was predicted and did not occur. The one factor ignored in the computation of the CSI is W, the number of times the event did not occur and was not predicted to occur. This often represents all but a small fraction of the predictions being tested.

		PREDICTION	
		YES	NO
EVENT	YES	X	y
	NO	Z	W

$$CSI = \frac{X}{X+Y+Z}$$

Figure 2. The Critical Success Index. X are successful predictions of the phenomenon, Y are failures to predict, and Z are false alarms. W are the occasions when the phenomenon was not predicted and did not occur

In the tests of Zeus at Seymour-Johnson, conducted during the fog season, W represented between 80 and 85 percent of the cases. A criticism of the CSI is sometimes made on the grounds that it "throws away data" because it ignores W. In actuality, it filters

the data, permitting one to observe comparatively slight differences between techniques. As shown in Figure 2, the CSI equals the number of hits (X) divided by the total number of pertinent cases (X + Y + Z). The CSI and skill scores for the three locations are listed in Table 1, based on statistics provided in GEOMET's final report.[2]

Table 1. Comparison of Human Forecasters at Seymour-Johnson AFB and Raleigh, North Carolina, With Zeus at Three Air Force Weather Detachments

		CSI	SKILL SCORE	N
HUMAN FORECASTERS	SEYMOUR-JOHNSON	.26	.23	27
	RALEIGH	.35	.24	40
ZEUS	DOVER	.22	.16	9
	FORT BRAGG	.00	.00	4
	SEYMOUR-JOHNSON	.40	.46	21
	COMBINED	.29	.35	34

3.1.1 RESULTS OF TESTING AT SEYMOUR-JOHNSON DURING 1988

The results presented in Table 2 were derived from two separate sources: real-time operation by Seymour-Johnson forecasters and operation by GL researchers using archived data. Only one forecast per day was considered, generally at 2000 hours local time, with a forecast period extending to the following morning. If fog occurred at any time between 2000 and 1000 hours the following morning, and if that fog reduced the visibility below 3 miles, then a reduced visibility event occurred for verification purposes. Some fog days fell into Category 3 because the lowest reported visibility was still in excess of 3 miles. The cases were all during the fall season, and archived data were used only when all data sources (or their equivalents) were available. Table 3 displays a breakdown of the test results, using the three visibility categories. In Table 2, all instances of Categories 1 and 2 are classified as "fog occurred," and all instances of Category 3 are considered "no fog" occurrences, whether or not fog was reported.

Earlier investigations of the fog forecast system, including GEOMET's final report,[2] demonstrated that the system underforecasts fog. This is in sharp contrast to the tendency

of human operational forecasters to overforecast the occurrence of fog, rationalizing that the penalty for failure to predict is much greater than the penalty for false alarms. It was decided to concentrate on occasions where the only type of error would be a failure to predict; therefore, only those days for which fog was reported the next morning were included.

As noted previously, it is not possible to duplicate operational conditions exactly

Table 2. Expert System Performer at Seymour-Johnson, Extended Field Proper

COMBINED STATS:

CAT 1 & 2 = FOG - YES
CAT 3 = FOG - NO

EVENT	PREDICTION YES	PREDICTION NO
YES	11	15
NO	1	3

TOTAL CASES =30

C.S.I. = 0.41

Table 3. Breakdown of Categories in Table 2

EVENT CATEGORY	PREDICTION 1	PREDICTION 2	PREDICTION 3
1	6	1	8
2	1	3	7
3	1	0	3

when using archived data. The inputs to the expert system that consist of intermediate forecasts, or judgment calls, are not archived. In running the system using archived data,

the archived observations and analyses were used, but whenever a forecast was required (for example, a response to a query concerning the location of the dominant high pressure area within the next 4 hours), the observational data for the latter time was referred to. This eliminated an intermediate source of error by assuring that all inputs were as accurate as possible. Therefore, one would expect that the results using archived data will always be better than those obtained in real time. However, in the present instance, there was no significant difference between the results obtained in real time and those using archived data. They have been combined in Table 2. The CSI obtained from these data is 0.41. Even allowing for the small sample size, this demonstrates that the prototype expert system can serve as a viable aid to the operational forecaster.

3.2 User Acceptance

As noted previously, the phrase "user acceptance" covers many aspects of the user-computer interface, and it both depends upon and affects the accuracy of the system. During the course of its field test of Zeus under its contract with GL, GEOMET collected comments from the users after each exercise of the system. The change in the manner in which observational data and the current synoptic situation are entered was in response to user comments. Additional changes designed to make the system less fallible are discussed in Section 5.

It was noted by GEOMET scientists that acceptance, in the sense of trusting the system's advice, increased with use. This went hand-in-hand with the high skill exhibited by the system, especially in those cases when the forecaster admitted that, without the expert system, he would have made an erroneous forecast. As the users became more familiar with the system, a phenomenon occurred that is quite common for small-to-medium expert systems.[4] The users reported that they began to anticipate the behavior of the system, and they were able to duplicate the logic and knowledge contained in the system without even turning the computer on. Because data input is still not automatic, the human forecasters were able to arrive at the same forecast as the computer in a shorter time. This is a perfectly acceptable result, indicating that the expert system operated first as an advisor and then as a teacher. A second, and less desirable, phenomenon results from the structure of the knowledge base, which does not have the flexibility usually associated with expert systems. There is no "fuzzy logic" built into the system. As a result, there are critical values of all inputs (for example, a dewpoint at sunset less than 39° F produces an automatic "no fog" forecast). Users who believe there will be fog during the night always enter a sunset dewpoint greater than 39° F, and they accept a forecast of no fog only for some reason other than too low a dewpoint.

---

4. Barr A. (1988) *The Future of Expert Systems*, Seminar, Lexington, Mass.

All instances that we can identify as cases where the user pushed the expert system toward a predetermined prediction was towards low visibility. There is no doubt that forecasters regard failures to predict low visibility as more serious errors than false alarms. Efforts to demonstrate this by using theCSI in its modified form and varying k were thwarted by the high number of failures to detect by both the human and the computer. Nevertheless, we can take it as given that forecasters are more willing to accept computer advisories that tend to be conservative rather than those having no bias at all.

Even more noteworthy when we consider the use of the expert system were the differences in the CSI at the different bases. Part of higher scores obtained at Seymour-Johnson may be attributed to the obvious fact that finer definitions of success were used at both Fort Bragg and Dover. Nevertheless, there does seem to be a higher level of success in the use of Zeus at Seymour-Johnson, a higher level that is not found when archived data were entered. Because there were no direct comparisons made between the human forecasters at Seymour-Johnson and at other locations, we cannot say whether or not the better forecasters also were better able to use the expert system, but we suspect this is so. The statistics gathered at Raleigh, North Carolina, using National Weather Service forecasters show that, in this case, the human forecasters were able to match the computer for the critical forecasts. This may indicate an upper limit on the accuracy of fog forecasts for this area unless a way is found to improve the computer guidance.

3.3 Expandability

At present, Zeus forecasts only radiation and advective fog as causes of low visibility. Deterioration in visibility resulting from frontal passage, or from blowing sand or snow, can easily be added to the system as separate rule-based modules. However, even at its present size, it is apparent that, for operational use, a faster method of data entry (data, in this instance, include the answers to questions posed by the computer regarding future synoptic situations) is imperative. The system cannot be expanded significantly without a redesign of the system architecture.

3.4 Transportability

Original specifications for the fog forecast system required that it could be run on the Z-100 computers at the base weather stations. The contractor immediately reported to us that no expert system shells can be used on the Z-100 because they were all written for IBM operating systems. Consequently, Zeus was developed on an IBM clone and tested on borrowed IBM-compatible computers. Further testing was done at Seymour-Johnson AFB using a Z-248. It was later learned[5] that the Z-100 can be made to emulate an IBM using

---

5. Roberts, D.K. (1988) Private communication.

software that mimics the IBM operating system. Zeus was operated successfully on a standalone Z-100 at GL, which differs from those issued to Air Force weather detachments only in having additional random access memory (RAM). It also lacks the IBM-compatible graphics package. Consequently, the map showing the sectors in which synoptic features may be located (part of Version 3 of the system) was not displayed. Attempts to run Zeus on versions of the Z-100 that do not have increased RAM failed.

It must be concluded, therefore, that even in its original configuration (without graphic interfaces), Zeus is not transportable to Air Weather Service (AWS) detachments with their present computer capability. This can be partially corrected by the addition of a memory board to the Z-100; however, the trend is toward increased use of graphic interfaces, not only in artificial intelligence (AI), but for other computer applications as well. The Z-248 now being obtained by the Air Force for certain locations can run the most recent version of Zeus, but probably would not be able to run future programs using light pens or other advanced interfaces.

4. ADAPTABILITY OF ZEUS BETWEEN SEYMOUR-JOHNSON AFB AND SHAW AFB

Zeus is a simple rule-based system using rules of thumb gleaned from interviews with forecasters at each of the three installations included in the study. Structure of the three versions of the system is identical: slightly over 200 IF-THEN rules arranged in three modules (synoptic, advective fog, and radiative fog) that are entered and exited in accordance with a data-driven decision tree. Differences between the three versions corresponding to the three locations are minimal and consist mainly in using different station reports to define the local weather regime. In testing the adaptability of the system, it was decided to transfer the rule base derived for Seymour-Johnson to Shaw, approximately 280 km to the southwest (Figure 1). Results of this effort were reported in Dyer.[6]

4.1 Analysis of the Knowledge Base

The first step in evaluating the system was to classify the rules according to type. The number of rules falling into each classification is less a function of the meteorological parameter being forecast than it is the result of the expert shell in which it was written, forecasting style of the expert or experts providing the knowledge base, thoroughness of the knowledge acquisition process, and the programming style of the person coding the knowledge base into the expert system. Zeus contains 10 rules that are strictly

---

6. Dyer, R.M. (1989) Adapting expert systems to multiple locations. *AI Applications in Natural Resource Management*, **3**(1): 11-16.

commonsensical. For example, several rules give the times of sunrise and sunset as functions of month followed by what may be classified as a "meta-rule":

IF the time is later than sunrise
AND the time is before sunset
THEN it is day
ELSE it is night

In addition to the 10 commonsense rules, there are 20 that are purely programming techniques, generally setting flags that permit a later rule to be expressed with a minimum number of clauses. Both commonsense and programming rules can be transferred to any location without modification.

Of the remaining 176 rules, 74 comprise the synoptic module. This module treats the presence, location, and movement of synoptic features, and relates them to future visibility. Many of these rules were expressed in combination with flags set by the programming rules mentioned previously. A translation of a typical synoptic rule follows:

IF a high pressure area is north-northeast of the station

AND the high pressure area is moving eastward

THEN there is potential for advective fog later

In transferring Zeus from Seymour-Johnson to Shaw, it was not necessary to change the synoptic rules, provided the locations of the synoptic features were always expressed relative to the station location. This is because the difference in location is relatively small: Both are East Coast stations, with the Atlantic Ocean to the east and the Appalachian Mountains due west. Transferring the system to a geograhic location with an entirely different pattern of moisture sources and topography would undoubtedly require modification of the synoptic module.

Thirty-two rules deal with the onset and dissipation of radiative fog. A typical example of such a rule is:

IF the visibility is less than 1 mile
AND a high pressure system is located east or south of the station
AND fog formation time was 10-14 hours after sunset
THEN radiative fog break time will be about 2 hours after sunrise

The exact way in which fog clearance time is determined by the critical location of the low pressure system and appropriate fog formation time were determined from climatological data at the site, and may differ at other locations; however, the general principle would remain: Radiative fog burns off after sunrise. The exact timing depends on fog depth, which, in turn, depends on factors such as the amount of radiational cooling. All of these factors can be determined from basic physical principles and are likely to vary only slightly from place to place.

Rules pertaining to advective fog are more site-specific than any other types of rules we have discussed thus far. A typical rule of this type is:

IF the month is April
AND the surface wind direction is >012
AND the surface wind direction is <102

THEN the surface wind direction is conducive to advective fog formation

There are 14 rules determining critical wind directions, including some that negate the possibility of fog because of downslope winds. In the original expert system development, these critical wind directions were obtained from climatological data and rules of thumb used at Seymour-Johnson. The adaptation of the system to Shaw was done by examining topographical maps, not by interviews with persnnel at Shaw AFB. Similarly, the rules calling for data and observations from locations surrounding Seymour-Johnson were modified by substituting stations in equivalent locations surrounding Shaw. Consequently, we can state that Zeus, written for Seymour-Johnson, can be adapted to Shaw without an extensive knowledge acquisition effort. The question then becomes: Does the system work as well at Shaw as it did at Seymour-Johnson?

4.2 Verification of the System at the Two Locations

The hypothesis that the modified version of Zeus performed adequately when transferred to Shaw was tested by comparing the CSI obtained at Shaw with that obtained at Seymour-Johnson. For the data and field tests at Seymour-Johnson, Zeus had a CSI of 0.38, while application of the modified version to archived data at Shaw yielded a CSI of 0.52. In both cases, the small sample size precluded any test of significance of the differences between CSIs. All we can say at this time is that using the CSI to compare verification statistics at the two locations, we note no deterioration in performance when Zeus was adapted from Seymour-Johnson to Shaw.

5. METHOD OF REDUCING THE FALLIBILITY OF THE PROGRAM

During the evaluation of Zeus, it became apparent that this system suffers from deficiencies in structure, logic, and user judgment calls (input). These deficiencies range from trivial to highly significant. Having said that, it should be emphasized that this effort was a proof-of-concept prototype, and, as such, it fulfilled its original objective. The comments below concern improvements that should be included in any future system intended for operational deployment.

5.1 Deficiencies in Logic

One of the primary benefits of expert systems is their ability to measure and make decisions on the amount and reliability of data leading to a final conclusion. Confidence in a forecast increases both with more inputs leading to the same conclusion and with the increasing certainty of each component input; however, in Zeus, there are only hard-coded confidence factors associated with each scenario end statement. We'd like to see, for example, narrow wind direction windows with higher certainty factors rather than more broadly inclusive ones. In addition, varying wind speeds should also have varying

confidence factors. These individual component certainties should combine to produce varying degrees of certainty for differing forecast outputs. The forecast would also be accompanied by a certainty or confidence factor, which would help the forecaster determine if an actual fog forecast should be made. Cutoff values of the certainty factor would be the determining force in this decision. Obviously, these cutoff values should be adjustable parameters to accommodate both learned experience as well as political biases (for example, local user authorities prefer false alarms over failures to predict). Defining these parameter certainties with pertinent value ranges is extremely difficult, because the body of knowledge is not already laid out in such a fashion. Finding appropriate certainty criteria might become a strenuous exercise in observe-and-adjust, ad infinitum. The ideal structure would be one that could digest, store, and make adjustments to parameters based on its data experience, as a human forecaster does. This falls under the AI discipline of "machine learning."

From our investigation, Zeus appears to be a difficult program to modify and maintain. Part of this difficulty lies in its not having the rule-base decision tree graphically written down, and part results from characteristics of the shell on which it was developed. While running Zeus, a user might be asked a question to which he may respond by typing in "Why?" Zeus then displays the rule generating the question, but it does not display the full thought chain incorporating this particular question. Future forecast systems to be deployed operationally should have the following capability: On demand, the program would present the currently tested conclusion as well as already affirmed data used in the path being tested. The intent of this capability is to be able to put the current question into the context of its usage, as opposed to just having the rule displayed. At the end of a run or as a separate function, one should be able to display or print the entire range of decision-tree possibilities. A capability of this sort would facilitate the inevitable alterations that will be required in an operational program.

Another problem in Zeus is the presence of apparent incomplete logic branches. For example, there are rules that are fired (become operative) only within time windows, such as "1300L < [TIME] < [SUNSET] + 4(hrs)"; however, other pieces of the puzzle (that is, times outside of that window) seem to be missing from the program. In particular, within the radiation module, there seems to be a cutoff around 2300L (current time of program run). In other words, if one enters a certain radiation fog scenario of data at a run time of 2000L, then the likely result will be a Category 1 fog forecast; however, if one enters 2330L or later for current time, then the same data produces neither a forecast category nor time. There was never any intended confinement of Zeus to certain hours of operation; therefore. this appears to be an error of omission. This situation results from the complexity of developing a large decision tree, scenario-type system such as Zeus. Each time one delineates a category or condition, then the complement of that condition must be considered. Specifically, if there is a conclusion that is predicated upon the current time being between

2000L and 2300L, then there must be conclusions that deal with the case of the current time being outside of that window.

The remainder of this section deals with specific rules in the expert system, or, in one case, what should be included as a rule. A number of the rules within Zeus use a quantity called "sunshine," which is basically the equivalent number of hours of sunshine experienced during the daytime before the night of the fog forecast (that is, one hour of no clouds equals one hour of sunshine). If this value is greater than 0.6 of the total possible number of hours of sunshine in the day, then fog will not be forecasted (radiation modules). Typically, this means that if there are more than 6 hours of sunshine in a day, then do not forecast fog for the next morning. This was repeatedly shown to be a cause of Zeus not forecasting actual fog occurrences. A study of climatological data often showed days with 10 hours of sunshine preceding a fog episode. Senior forecasters at Seymour-Johnson AFB expressed the opinion that if this parameter has validity as used, then it must also have a seasonal relationship that has not been included in Zeus. In other words, the parameter may be valid in a summer scenario, but not in a fall one. To accommodate this situation, Seymour-Johnson forecasters would frequently input bogus values less than 6 for "sunshine." Unfortunately, the hours of sunshine are also used in calculating fog formation times, so any bogus values will result in contaminated forecast times. Our immediate recommendation is to drop this particular rule and use the sunshine parameter only in the formation-time calculation. At this point, we are undecided as to how or if this factor should be utilized in future efforts.

Another set of rules deals with the dewpoint depression (that is, temperature minus dewpoint temperature). Using current run-time values of temperature and dewpoint temperature for specific local stations, Zeus calculates the average dewpoint depression. When calculated, this value is compared to a fixed value of 10. If it is less than or equal to 10, then fog is possible. Ten is used regardless of the time of observations. One would logically expect the dewpoint depression to be large around sunset but decrease throughout the night. In other words, a 10° spread at sunset is quite different from a 10° spread at 0200. We'd like a dewpoint spread vs time curve determined in conjunction with cloud conditions to be used in this kind of expert system. The question we'd like the expert system to address is, "Given current conditions, will the dewpoint spread decrease enough by the morning to form fog?" This is another place where certainty factors should be assigned.

Forecasters report a dissatisfaction with the current rule that asks the user if there was rain in the afternoon. The intent of the rule determining a sufficient moisture source is obvious and unchallenged, yet it omits the frequent case of the ground being soaked from earlier rains. Forecasters at Seymour-Johnson spoke of wet periods that left the ground soaked for several days following the final rain. The more appropriate question for Zeus is to ask whether or not the local terrain is <u>unusually</u> wet from recent rainfall. Admittedly, this required a judgment call by the user. If this is to be avoided, then some record of rainfall rate

over the past week must be requested or maintained in the computer's data base, and a cutoff value for the program established.

A similar comment deals with the reverse case, that is, unusually dry conditions. There is not a rule in Zeu dealing with this issue. Even with all normal input fators being conductive to fog formation, fog will probably not form if the local terrain is extremely dry (moisture sponge). This should be incorporated in the knowledge base in some manner.

Feedback from Seymour-Johnson forecasters also indicated the absence from Zeus of one of their most reliable precursors of fog, anomalous propagation (AP) on their radar. AP indicates that the needed low level inversion is present, thereby capping the available moisture. In practice, this factor dramatically increases the forecaster's confidence that fog formation will occur in the next several hours. In Zeus, there is no convenient mechanism for inserting new data sources. As expressed elsewhere in this report, one of our prime hopes for a future operational system will be the capability for local forecasters to tailor a fundamental and generic system to their specific site. To do so will require the ability of adding new and specific data sources to the already existing base program. Work on this concept is now in progress.

5.2 Deficiencies in Structure

The first and foremost structure complaint made by the users is the pervasive use of programming flags that, by their nature, are meaningless to the user. Examples of this are "Condition-one is met" and "nochange-h2 flag is set." When the user inspects the fired rules, these statements offer no insight into the reasoning process intended, and, in fact, act to confuse him. GEOMET attempted to mitigate this characteristic with rule notes that are displayed along with the rule. This is, however, not sufficient to satisfy a user's need to understand the specific thought process used by the program. These flags also make the job of maintaining and updating the knowledge base extremely difficult. To demonstrate the seriousness of this problem, we present the following rule in Zeus:

Rule Number 70

IF: [TIME] > [SUNSET]
AND [TIME] < 1200
and Dummy-1a is set
and AF intensity flag is set
and Flag special is set to Nochange-H2 or Nochange-H2F or
Nochange-H2S

THEN: (TEXT for printout) "Assuming maintenance of Atlantic flow and clear skies to coast, then visibility should deteriorate this coming evening. Keep checking surface observations out to the coast."

and Dummy-50 is not set

ELSE: Dummy-50 is set

NOTE: WHY.... THIS IS RELATED TO THE CHANGING SYNOPTIC SITUATION, WHERE A HIGH IS GOING TO SET-UP AND STALL IN MID-ATLANTIC STATES.

CAUTIONARY MESSAGE ALERTS FORECASTER TO THE POTENTIALS OF THIS SITUATION.

REFERENCE: INTERVIEWS

Hopefully, the reader finds the cryptic nature of this rule disturbing. The text statement as well as the programming note at the bottom are ineffective in clarifying the workings of this rule. Let's pursue this process further as if we were operationally trying to investigate the thought process behind the rule. First, we must determine what the condition "AF intensity flag is set" means. We search through the 207 rules to find the rule defining it, which happens to be rule number 27 in this case.

RULE NUMBER 27

IF: [TIME] > 1000
and [TIME] < 2400
and [FWIND12S] >= 10
and [FWIND12S] <= 30
and [FWIND12D] > 040
and [FWIND12D] <= 210
and [CSCTD SUNSET] > 39
and Dummy-1a is set
and Sky, from GSB to coast, is CLR to SCT

THEN: AF intensity flag is not set

ELSE: AF intensity flag is set

NOTE: WHY.... CHECKS FOR PROPER CONDITIONS GIVEN ATLANTIC FLOW THAT WILL ALLOW FOR ADVECTIVE FOG. THIS IS A MASTER RULE THAT RUNS BETWEEN 1200 NOON AND MIDNIGHT, AND FOR CURRENT VIS GREATER THAN 3 MILES.

REFERENCE. INTERVIEWS

Note that the condition "AF intensity flag is set" is treated in the else part of rule 27. For this to be true, at least one of the condition statements in rule 27 must be false, yet as many as all of them may be false. Most likely, it will not be obvious to the user which are and which are not true. Also note that, within the IF condition statements, there is another cryptic flag, "Dummy-1a is set." Therefore, we must investigate the rule defining this flag. In some cases, this nesting of flags proceeds four and five rules deep. Hopefully, from this example, the reader can sense the frustration and bewilderment that a real-time user would feel if he pursued the reasoning process of Zeus. We have inspeted the program in a much more deliberate fashion than afforded to the typical user, yet we still are very confused by much of the rule structure. In view of this, we recommend a structure that shows nothing but sensible text or equations.

Automatic input of data should be considered imperative in future weather expert systems. The amount of time spent by the forecaster in entering all the input data is

excessive and could be better utilized. Even after the program is completely loaded, which in itself is over a 1-minute process, the user requires a substantial amount of time to load the inital input data. From our observations, 20 minutes seems to be a good representative value for this input time. GEOMET noted this situation in its report.[2] The manual form of data entry was reluctantly adopted, because the costs and complications involved in an automated system were prohibitive considering the prototype scope of this contract. Future efforts should use the automated approach. Ideally, a weather data base should be maintained continually and tapped when needed by the expert system.

Within Zeus, there is an interesting division of labor that unfortunately results in occasional incompleteness of output. Specifically, there are two separate forecasts going on in the program, forecast time and forecast category. The intent is that the two link up for a final coherent forecast; however, at times, Zeus puts out a forecast time without a forecast category or with a Category 3 (good visibility) forecast. The result is forecaster confusion as well as skepticism about the expert system. Currently, in the structure of the rule base, the fog formation calculation helps drive the category forecast. In other words, a typical scenario end statement has as one of its conditions "Radfoghour calculation is done." It would seem more appropriate to forecast a category of fog and then trigger an equation or module to make the formation time forecast.

Zeus requests current observation data (temperature and dewpoint temperature) from the user for a fixed set of stations. These data are used to calculate the average temperature dewpoint difference (dewpoint spread), which is then used in the rule base. A couple of facts make this format undesirable. First, several of the stations are not 24 hours per day, 7 days per week operations, and, as such, frequently are not available to supply the needed input. Forecasters are then forced to estimate values using available data. Secondly, forecasters don't always need or want to consider the dewpoint spread at all surrounding stations. Instead, they are more interested in a more narrowly defined window of upwind stations. We'd like to see a system that coordinates this data request with trajectories derived perhaps from geostrophic analysis (to neglect small-scale surface aberrations). We envision a sliding data window (say, triangular) that is automaticaly adjusted with the trajectory inputs. In this way, there would be a more realistic fine-tuning of input data.

5.3 Problem of Judgment Calls

A number of problems arise in the evaluation of Zeus because of "judgment call" questions. The prime example of this is in entering synoptic feature data. In other words, Zeus will ask the user what "quadrant" (there are six quadrants) the high is in. Unfortunately, analysis (pressure or height contours) maps are not as neat and simplistic as Zeus implies in its questions. Much confusion was expressd by the users in interpreting what scale of feature to consider. Often, there will be small-scale (meso- or micro-) features embedded in

an area of contrasting charcteristic. For example, a meso-low can be in the middle of a broad-scale ridge or high. Also, features such as lows and highs can be ridges or valleys that extend through different quadrants and occupy more than one quadrant simultaneously. Other judgment calls include predicted movement and intensity of cold and warm fronts, as well as expected rainfall. In an ideal system, we'd like these judgment calls to be answered by the program using the available numerical products. If the user wanted to change an entry, he would call up the parameters used, make his change, and then rerun the program as is done in the current system.

Eliminating all user judgment calls was far beyond the intended scope of Zeus. However, as a future expectation, we enthusiastically endorse the concept of the "behind-the-scenes" expert system making these judgment calls directly from data sources. We say this with the presumption that any aspect of the decision process can be queried and/or altered by the human user.

6. DEVELOPING ADAPTABLE SYSTEMS

The exercise described in Section 4, in which an already existing system for one location was adapted for use at another, demonstrated that an entirely different approach must be taken to the development of expert systems designed to be used worldwide for meteorological purposes. The new approach, described here, is to start with basic physical and meteorological principles, gradually adding on layers of observations and specific *applications of generic knowledge*, until the most superficial layer is reached--that of the site-specific local rules of thumb. Figure 3 illustrates how the knowledge acquisition might be divided between the developer of the system and the user. The first level of the system consists of all pertinent physical and meteorological principles. Taking the fog forecast system as an example, the fact that when air is cooled to its dewpoint, condensation occurs, would be found in this level of the system. The second level of the knowledge base might contain the information that, for coastal areas, an appropriately located ridge of high pressure might sustain oceanic flow, and hence favor continued advective fog. Both these levels are part of the deep knowledge base to be supplied to all users of the system. All users would make use of the first layer; those interested in forecasting only for continental locations could ignore the particular example of second level knowledge cited here. It is possible that the second level knowledge base would be contained in modules described in general terms (tropical island, temperate coast, sub-polar continent, etc.), and the users would install only those modules pertaining to the stations for which they wish to forecast. advective fog. Both these levels are part of the deep knowledge base to be supplied to all users of the system. All users would make use of the first layer; those interested in forecasting only for continental locations could ignore the particular example of second level knowledge cited here. It is possible that the second level knowledge base would be contained in modules described in general terms (tropical island, temperate coast, sub-polar continent, etc.), and the users would install only those modules pertaining to the stations for

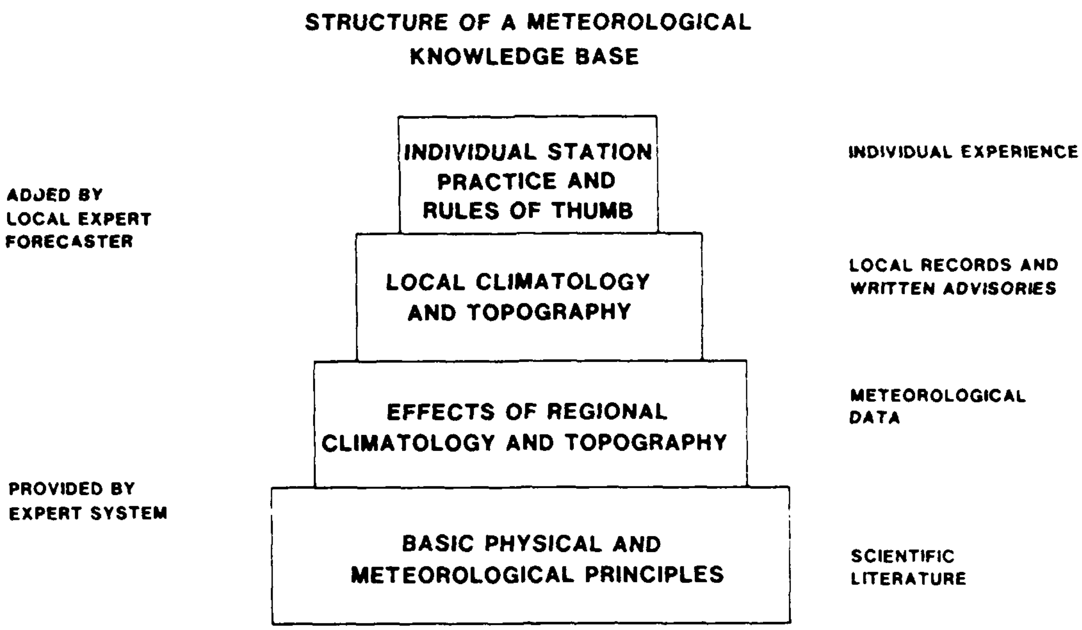

Figure 3. Levels of Knowledge in a Weather Forecasting System

which they wish to forecast. The third and fourth levels of the knowledge base require site-specific input. In the third level, for example, the locations of sources of moisture and of topographic features would be entered. This would be used by the system to generate the critical wind directions that would allow (or prohibit) advective fog formation. The final, most superficial level contains the local rules of thumb and the sources of data currently in use at the individual station. The Seymour-Johnson version of this level would contain, for example, the names of the coastal stations whose dewpoint depressions are averaged to obtain a measure of the moisture content of the incoming air mass; the fact that 10° F is considered a critical value of this quantity; and that unless the dewpoint is at least 39° F, radiation fog is unlikely to occur at Seymour-Johnson. This final level is the one most subject to change as data sources are added or phased out, or as new on-duty forecasters discover better rules of thumb. It is conceivable that an Air Force meteorologist might be called on to forecast for a location at which there is no past experience and for which local records are scanty. Under these circumstances, the system would be used initially with only the first two levels of knowledge.

6.1 Deep Knowledge in Fog Prediction

Some examples of factors causing or preventing fog are listed in Table 4.[7] The fog-causing factors can be categorized as either evaporation or cooling. Fog-dissipating

---

7. Pettersen, S. (1956) *Weather Analysis and Forecasting*, McGraw-Hill, AWS Boston, Mass., 2nd Edition.

Table 4. Physical Processes Causing Fog Formation and Dissipation. These would be incorporated into the "deep knowledge" of a generic fog forecast system, corresponding to the bottom layer of Figure 3

1. Change of State

   Evaporation from rain or a water surface (such as a river, lake, or ocean) that is warmer than the air causes fog.

   Sublimation or condensation on snow is a fog-dissipating process.

2. Change of Temperature

   Warm, moist air may be cooled to its condensation level by:

   a. adiabatic upslope motion;
   b. motion across isobars toward lower pressure;
   c. falling pressure;
   d. radiative cooling
   e. advection of warmer air over a colder surface.

   Fog will dissipate when the air is heated to a temperature above its dewpoint by:

   a. adiabatic upslope motion;
   b. motion across isobars toward higher pressure;
   c. rising pressure.

3. Mixing

   Vertical mixing is an important factor in dissipating fogs and forming stratus clouds.

processes are their opposites: condensation (or sublimation) and heating, with vertical mixing an additional factor in transforming ground fog to stratus clouds. Rules incorporating the factors listed in Table 4 would constitute the deepest level of the knowledge base. The next level, topographical and climatological factors, would be of the type listed in Table 5. These are the two types of knowledge depicted in Figure 3 that would be incorporated into the expert system delivered to the users.

6.2 Local Factors and Rules of Thumb in Fog Prediction

The next layer of the knowledge base would contain information about the local topography and climatology for the location for which the predictions are being made. Examples of this might be the rule at Seymour-Johnson that if the dewpoint at sunset is less than 39° F, then overnight radiation fog will not occur. This information is similar to that contained in the forecaster worksheets and advisories for newly arrived forecasters. The final level contains the rules of thumb used by a particular expert forecaster. An example of that might be the forecaster arriving for the eary morning shift noting whether there was dew on the grass or condensation on automobile windshields, and using that information as evidence of saturation levels.

Table 5. Some Examples of General Regional and Topographical Factors Influencing Fog Formation. These, together with regional climatology, would be incorporated into the next-to-bottom layer of a generic fog forecast system

Sea fogs are frequent in the vicinity of cold ocean currents.
Advective fog formation along coasts of oceans of other large bodies of water occurs most frequently in local spring, when the water is colder than the surrounding land.
Radiation fog occurs most frequently on cold cloudless nights. It cannot form if a heavy overcast prevents radiation cooling. It is also more sensitive to wind and shallower than advective fogs.
Fog forms in valleys when moist air becomes trapped under stagnant anticyclonic conditions.
Because of diurnal heating, land fog is more likely to occur in early morning and least likely in the afternoon.
Fog rarely develops in the interior of continents, even when snow-covered.
Upslope fogs are frequent on the windward side of mountain slopes They are rare at low elevations.
The larger the difference between the temperature of the air and that of the underlying surface, the stronger the wind in the presence of which an advective fog can be maintained.
The shallower the fog, the more quickly it dissolves as a result of diurnal heating.

7. CONCLUSIONS

The experience in developing and evaluating the expert system for forecasting fog has demonstrated that expert system technology is not only feasible and desirable, but it is an inevitable technology in operational weather forecast stations of the future. It is also apparent that the knowledge base must be expressed in the layered structure described in Section 6. A third conclusion is that there will have to be developed a domain-specific shell for meteorological analysis and prediction. This implies that the shell will contain information of the characteristics of meteorological parameters and their interrelations. For example, a meteorological shell will accept dewpoint relative humidity or absolute humidity as measures of moisture content; will be able to convert systems of units; and will substitute derived or default values when observations are not available. The combination of a layered knowledge base and domain-specific meteorological shells will permit the implementation of weather forecast expert systems at all Air Force installations.

References

1. Stunder, M., Dyer, R., and Koch, R. (1987) The use of an expert system in assisting forecasters in visibility predictions, *3rd Conference on Interactive and Processing Systems in Meteorology, Oceanography, and Hydrology*, pp. 5206-5207.

2. Stunder, M., Koch, R., Sletten, T., and Lee, S. (1987) *ZEUS: A Knowledge-Based Expert System that Assists in Predicting Visibility at Airbases*. AFGL-TR-87-0019, AD A184197.

3. Donaldson, R., Dyer, R., and Kraus, M. (1975) An evaluator of techniques for predicting severe weather events, *Ninth Conference on Severe Local Storms*, pp. 321-326.

4. Barr, A. (1988) *The Future of Expert Systems*, Seminar, Lexington, Mass.

5. Roberts, D.K. (1988) Private communication.

6. Dyer, R.M. (1989) Adapting expert systems to multiple locations. *AI Applications in Natural Resource Management*, **3**(1): 11-16.

7. Pettersen, S. (1956) *Weather Analysis and Forecasting*, McGraw-Hill, AWS Boston, Mass., 2nd Edition.